

Deconvolution, convex optimization, non-parametric empirical Bayes and treatment of non-response.

Eitan Greenshtein

Israel Central Bureau of Statistics; e-mail: eitan.greenshtein@gmail.com

and

Theodor Itskov

Israel Central Bureau of Statistics ; e-mail: itsmatis@gmail.com

Abstract.

Let (Y_i, θ_i) , $i = 1, \dots, n$, be independent random vectors distributed like $(Y, \theta) \sim G^*$, where the marginal distribution of θ is completely unknown, and the conditional distribution of Y conditional on θ is known. It is desired to estimate the marginal distribution of θ under G^* , as well as functionals of the form $E_{G^*} h(Y, \theta)$ for a given h , based on the observed Y_1, \dots, Y_n .

In this paper we suggest a deconvolution method for the above estimation problems and discuss some of its applications in Empirical Bayes analysis. The method involves a quadratic programming step, which is an elaboration on the formulation and technique in Efron(2013). It is computationally efficient and may handle large data sets, where the popular method, of deconvolution using EM-algorithm, is impractical.

The main application that we study is treatment of non-response. Our approach is nonstandard and does not involve missing at random type of assumptions. The method is demonstrated in simulations, as well as in an analysis of a real data set from the Labor force survey in Israel. Other applications including estimation of the risk, and estimation of False Discovery Rates, are also discussed.

We also present a method, that involves convex optimization, for constructing confidence intervals for $E_{G^*} h$, under the above setup.

1. Introduction, Preliminaries and Examples.

Consider a general empirical Bayes setup, where (Y_i, θ_i) , are i.i.d., $i = 1, \dots, n$, distributed like $(Y, \theta) \sim G^*$, and the conditional distribution of Y conditional on θ is F_θ , $\theta \in \Theta$. The marginal distribution of θ under G^* is denoted G . Suppose we only observe Y_1, \dots, Y_n , and we should estimate the parameters $\theta_1, \dots, \theta_n$. It is often the case that G is unknown and should be estimated in the process of estimating the unknown parameters. We concentrate on the non-parametric

empirical Bayes setup where G is completely unknown, as opposed to the parametric setup where G is assumed to be a member of some parametric family.

We have two main novel contributions in this paper. One is suggesting a new deconvolution method, for the purpose of estimating G by a corresponding estimator \hat{G} . The deconvolution method is based on quadratic programming. Note, an estimator \hat{G} for G induces a corresponding estimator \hat{G}^* for G^* , through $d\hat{G}^*(y, s) = dF_s(y)d\hat{G}(s) \equiv dG^*(y|\theta = s)d\hat{G}(s)$. The other main contribution is a nonstandard application of deconvolution and empirical Bayes to the problem of treating non-response. Other applications are also described.

In the canonical examples of empirical Bayes the ultimate goal is the estimation of the parameters θ_i , $i = 1, \dots, n$, based on the observed Y_1, \dots, Y_n . However, our main interest and emphasis is on estimating various functionals of the form $E_G h$ and $E_{G^*} h$ for various functions h . We also consider the more general setup where (X_i, Y_i, θ_i) , $i = 1, \dots, n$ are i.i.d., distributed like $(X, Y, \theta) \sim G^*$; the joint distribution, G , of X and θ is completely unknown, while $G^*(Y|X, \theta)$ the conditional distribution of Y conditional of X and θ is known. We observe n independent realizations $\mathcal{T}(X_i, Y_i)$, $i = 1, \dots, n$ for some function \mathcal{T} . Here the pair (X_i, θ_i) may be considered as the ‘parameter’ that determines the conditional distribution of Y , however, unlike the former setup, the ‘X-part’ of the ‘parameter’ is observed through $\mathcal{T}(Y_i, X_i)$, i.e., the parameter (X, θ) is not completely latent. The goal is again to estimate $E_{G^*} h(X, Y, \theta)$ for various functions h . The estimators are of the form $E_{\hat{G}^*} h$, for a ‘deconvolution-estimator’, $d\hat{G}^*(x, y, s) = dG^*(y|\theta = s, X = x)d\hat{G}(x, s)$.

In Section 3 we present a method for constructing a confidence interval for quantities of the form $E_{G^*} h(X, Y, \theta)$, based on $\mathcal{T}(X_i, Y_i)$ $i = 1, \dots, n$. The main idea of that method is defining an appropriate convex optimization problem, where the target function is linear and the constraints are convex.

In the rest of this section we elaborate on a few empirical Bayes examples, where it is desired to estimate quantities of the form $E_{G^*} h$. Our primary example is the problem of treating non-response. Finally, we explain why estimation of G by \hat{G} and then estimation of $E_G h$ by $E_{\hat{G}} h$ is a good alternative to, say, mle estimation of $\frac{1}{n} \sum h(\theta_i)$ by $\frac{1}{n} \sum h(\hat{\theta}_i)$, where $\hat{\theta}_i = \hat{\theta}_i(Y_i)$ is the point-wise mle estimator of θ_i , $i = 1, \dots, n$.

1.1. Examples.

Deconvolution, Empirical-Bayes, and Estimation of the risk.

In the canonical examples of Empirical Bayes, the ultimate goal is to estimate the individual parameters θ_i , $i = 1, \dots, n$. In such problems the estimation of $E_{G^*} h$ for various h could still be central, as demonstrated in the following.

Let $\delta(Y)$ be a decision function and $L(\theta, \delta(Y))$ a loss function. Of a primary

interest is the quantity

$$E_{G^*}L(\theta, \delta(Y)) = E_G R(\theta, \delta) = E_G h_\delta(\theta); \quad (1)$$

here $Y \sim F_\theta$, and $R(\theta, \delta) \equiv h_\delta(\theta)$ is the risk of δ conditional on θ . Thus, the quantity in (1) is the Bayes risk that corresponds to the decision function δ , under the loss L and the prior G . The Bayes procedure is

$$\delta^B = \operatorname{argmin}_\delta E_G h_\delta(\theta).$$

Uniformly good estimation of $E_G h_\delta(\theta)$ over all δ , yields good estimates of δ^B .

Once an estimator \hat{G} for G is obtained, a natural approach is to let

$$\hat{\delta}^B = \operatorname{argmin}_\delta E_{\hat{G}} h_\delta(\theta). \quad (2)$$

Under a squared loss, the estimated decision function in (2) is

$$\hat{\delta}^B(y) = E_{\hat{G}^*}(\theta | Y = y).$$

More generally, under squared loss, in the setup where $(X_i, Y_i, \theta_i) \sim G^*$ and we observe $\mathcal{T}(X_i, Y_i)$, $i = 1, \dots, n$, a natural estimator for θ based on $\mathcal{T}(X, Y)$ is:

$$\hat{\delta}(\mathcal{T}(X, Y)) = E_{\hat{G}^*}(\theta | \mathcal{T}(X, Y)).$$

The case where X and Y are independent conditional on θ is of a special interest, e.g., as in our simulation section.

We should remark that \mathcal{T} may be a randomized transformation. In Brown et.al. (2013), the set up is $(X_i, Y_i, \theta_i) \sim G^*$, $i = 1, \dots, n$, are i.i.d, where the conditional distribution of Y conditional on θ is $Poisson(\theta)$, while $X \sim Poisson(h)$ is independent of θ and Y . In that paper even though a direct observation of Y is available, the approach is to estimate the optimal decision function with respect to the artificially "corrupted" observation $\mathcal{T}(X, Y) = X + Y$, with $h = h_n \rightarrow 0$, as $n \rightarrow \infty$. This approach is shown to have advantages relative to, say, the classical EB estimator for a Poisson parameter, suggested by Robbins. In the sequel we will not consider randomized \mathcal{T} , although it is within our formulation.

There are common examples, e.g., Poisson, Normal, where an estimator for δ^B may be obtained directly without the estimation of G and application of (2). On the direct approach for the estimation of δ^B , versus approaches that involve the estimation of G , see, e.g., Efron (2013). On direct estimation of δ^B in the normal case see, e.g., Brown and Greenshtein (2009); on direct estimation in the Poisson case see, e.g., Brown, et.al. (2013).

Deconvolution and Variations on False Discovery Rate

Problems that involve estimation of $E_{G^*}h$ are related also to the problem of estimating false discovery rates (FDR), see, Benjamini and Hochberg (1995).

Let $(Y_i, \theta_i) \sim G^*$, $i = 1, \dots, n$, be independent where conditional on θ_i , $Y_i \sim F_{\theta_i}$, $i = 1, \dots, n$. Consider first the problem where it is desired to estimate

the proportion of indices i for which $\theta_i > C$. When n is large and G , the marginal of θ is known, a reasonable trivial estimator, is $P_G(\theta > C)$. Note that, $P_G(\theta > C) = E_G h(\theta)$ for the function h which is the indicator of the event $\{\theta > C\}$. When G is unknown and estimated by \hat{G} , the induced estimator is $E_{\hat{G}} h$.

We now treat the more general (FDR) problem. In order to fix ideas consider the case $F_{\theta_i} = N(\theta_i, 1)$. Suppose that we consider observations i with $Y_i > A$ for some A as "suspected discoveries", while we consider as "true discoveries", observations for which $\theta_i > C$. In order to estimate the proportion of true discoveries among suspected discoveries we should estimate the quantity $E_{G^*} h$ for the function h which is the indicator of the event $((\theta_i > C) \cap (Y_i > A))$.

When it is desired to estimate the proportion of "true discoveries" among suspected discoveries for a given realization, the following perspective and alternative approach might be beneficial. Let G^{*t} be the conditional distribution of (Y, θ) conditional on the event $Y > A$, we treat the observations (Y_i, θ_i) with $Y_i \leq A$, as truncated and the remaining ones are treated as i.i.d., observations from G^{*t} , where $G^{*t}(y|\theta) = F_\theta(y|Y > A)$. Let G^t be the marginal distribution of θ under G^{*t} and \hat{G}^t its "deconvolution-estimate", let h be an indicator of the event $\theta > C$, we may estimate $E_{G^t} h$ by $E_{\hat{G}^t} h$. See Greenshtein et. al. (2008), for treatment of a related problem.

Deconvolution and Treatment of Non-Response. A main novel contribution of this paper, is an application of our deconvolution method to treat non-response. The proposed treatment of non-response does not involve the, often assumed and seldom verifiable, assumption of Missing At Random (MAR), conditional on some covariates.

Let $\mathbf{S} = \{i_1, \dots, i_n\}$ be a random set of indices that correspond to randomly sampled items from a finite population of size N , indexed by $\{1, \dots, N\}$. Those are the indices of the items in the population who i) were randomly sampled for a survey ii) responded.

Suppose, it is desired to estimate the total $T = \sum_{i=1}^N X_i$ in the population, based on the n available observations. Let I_i be an indicator of the event "item $i \in \mathbf{S}$ ", $i = 1, \dots, N$. Let $p_i = E(I_i)$. Then

$$\hat{T} = \sum_{i \in \mathbf{S}} \frac{X_i}{p_i},$$

is the Horvitz Thompson estimator for T . It is an unbiased estimator, as may be seen from the representation

$$\hat{T} = \sum_{i=1}^N \frac{X_i}{p_i} I_i. \quad (3)$$

Typically, p_i , $i = 1, \dots, N$ are unknown although the sampling probabilities are known. This is since the corresponding response probabilities are unknown. Thus, the above estimator can not be applied.

We will approximate (3), by a nonparametric empirical Bayes modeling together with a deconvolution step. Consider a situation where there is an additional covariate Y_i for every item i , $i \in \mathbf{S}$, such that $Y_i \sim F_{p_i}$. In one example, that we will give, Y_i is the number of visits until a response was obtained; in another example Y_i is the number of responses of item i in a longitudinal panel survey, where each sampled item is attempted to be interviewed four times in four consecutive months.

We model the observations (X_i, Y_i, p_i) , $i \in \mathbf{S}$ as i.i.d $(X_i, Y_i, p_i) \sim G^{*t}$. Here G^{*t} is the conditional distribution of (X_i, Y_i, p_i) conditional on $i \in \mathbf{S}$, which is different than G^* the distribution of (X_i, Y_i, p_i) , $i = 1, \dots, N$. A natural estimator for (3) is $nE_{G^{*t}} \frac{X}{p} \equiv nE_{G^{*t}} h(X, p)$, for $h(X, p) = \frac{X}{p}$. This treatment is under truncation, where we have no knowledge about the observations that correspond to indices i , $i \notin \mathbf{S}$. Under censoring, when there exists partial information about items with index i , $i \notin \mathbf{S}$, related ideas will be applied. The formal distinction and different treatment under truncation versus censoring will be explained and demonstrated in Sections 4,5.

A general reference to sampling is, e.g., Lohr(2009). A reference for missing data and non-response issues is, e.g., Little and Rubin (2002).

Non parametric maximum likelihood estimation of G

The first study of the estimation of G , under the above setup, was conducted by Kiefer and Wolfowitz (1956). They suggested to find the non-parametric mle for G , and also gave conditions under which the non-parametric mle estimator \hat{G} converges weakly to the true G . Estimation of $E_G h$ by $E_{\hat{G}} h$ is often much better than estimating the individual parameters θ_i , say by an mle, $\hat{\theta}_i(Y_i)$, and then average, to obtain the estimator $\sum^n h(\hat{\theta}_i)/n$. This is demonstrated in the following Example 1.

Example 1. Consider the Normal example where $F_{\theta_i} = N(\theta_i, 1)$. Let h be the function $h(\theta) = 1/\theta$. Suppose it is known that the support of G is bounded below by 0.5, but otherwise it is completely unknown. Then the mle for θ_i is $\hat{\theta}_i = \max(0.5, Y_i)$. Now suppose that the true G has a point mass at 1. By Kiefer and Wolfowitz (1956), the mle \hat{G} for G converges weakly to G , so $E_{\hat{G}} h \rightarrow 1$. However, a quick simulation shows that $\frac{1}{n} \sum \frac{1}{\hat{\theta}_i}$ converges to 1.19.

Estimation of sums of the form $\sum_i h(Y_i, \theta_i)$, was studied by Zhang (2005). Further examples may be found there, as well as a study of the efficiency of certain estimators.

The rest of the paper is organized as follows. In Section 2, our deconvolution method is explained. In Section 3, we present a method that involves convex optimization, to construct confidence intervals for quantities of the form $E_{G^*} h$. In Section 4 we present ‘empirical Bayes type Horvitz Thompson’ estimators in

the context of treating non-response. In Section 5, the derivation and performance of those estimators are illustrated through a simulated practical example. In Section 6, we demonstrate our method for treating non-response, through a real data set from the Labor Force Survey in Israel.

2. Deconvolution using quadratic programming.

In this section we present a deconvolution algorithm which involves quadratic programming.

Our deconvolution is a method for deriving a Non-Parametric Maximum Likelihood Estimator (NPMLE) for a ‘prior’ G . We use the term deconvolution in a wide sense, that includes identifying mixtures, as studied, e.g., by Lindsay (1995), Lindsay and Roeder (1993), Lee et.al., (2013), and, of course, the fore mentioned seminal paper of Kiefer and Wolfowitz (1956). Our quadratic programming approach, rather than the more common EM-algorithm, is in line with the general suggestion and advocacy of Koenker and Mizera (2013) for the usage of convex optimization. It may be applied on high dimensional problems with tens of thousands of observations and general mixing G , where the complexity of EM algorithms makes them impractical.

In the following subsection we treat the problem of estimating the marginal distribution of a latent-variable/unknown-parameter. The same ideas apply to the more general problem of estimating the joint distribution of a latent variable and an observed variable. We present the ideas in two stages where the general case is formulated in subsection (2.2).

Our approach is based on the setup and formulation in Efron (2013). We elaborate more by defining and solving an appropriate quadratic programming problem.

2.1. Deconvolution for the estimation of the marginal distribution of a latent variable.

Consider a standard empirical Bayes setup, as described in the introduction, where $(Y_i, \theta_i) \sim G^*$, are i.i.d., $i = 1, \dots, n$. We assume discrete distributions, in particular F_θ , $\theta \in \Theta$, are discrete with a common finite support denoted $\{y_1, \dots, y_J\}$, and G is discrete with a given support $\{s_1, \dots, s_K\}$. The treatment of the continuous cases may be done through discretization. In principle the discretization of the Y -variables should be more delicate as the number of observations increases, but the ‘right’ way of discretization is beyond the scope of this paper. Our main examples and applications in sections 4-6 involve discrete observations Y_i , $i = 1, \dots, n$, specifically, censored Geometric and Binomial. The considerations that are involved in the discretization of G have to do with the complexity of the estimation algorithm.

Our observations Y_i , $i = 1, \dots, n$, are independent and identically distributed like a random variable Y . Denote their discrete density by $\mathbf{f} = (f_1, \dots, f_J)'$,

where $f_j = P(Y = y_j)$, $j = 1, \dots, J$. Denote

$$p_{jk} = P(Y = y_j | \theta = s_k), j = 1, \dots, J; k = 1, \dots, K.$$

Denote the density of the discrete distribution G by $\mathbf{g} = (g_1, \dots, g_K)'$, where $g_k = P_G(\theta = s_k)$, $k = 1, \dots, K$. Denote by P the $J \times K$ matrix $P = (p_{jk})$.

Then:

$$\mathbf{f} = P\mathbf{g}. \quad (4)$$

Recall, the support of G is known (or practically approximated by a dense grid $\{s_1, \dots, s_K\}$), it is the density \mathbf{g} that should be estimated. We now reduce the problem through sufficiency. Note that a sufficient statistic is $\hat{\mathbf{f}} = (\hat{f}_1, \dots, \hat{f}_J)'$, where \hat{f}_j is the proportion of observations among Y_1, \dots, Y_n , that had the value y_j , $j = 1, \dots, J$. Now, $\hat{\mathbf{f}}$ is a scaled multinomial vector with mean \mathbf{f} and a corresponding covariance matrix $\Sigma_{\hat{\mathbf{f}}}/n$. Its distribution is asymptotically multivariate normal. Note, that there is a linear dependence, thus the corresponding covariance matrix $\Sigma_{\hat{\mathbf{f}}}^{-1}$ does not exist. We may replace $\hat{\mathbf{f}}$ by the sufficient statistic $\hat{\mathbf{f}}^* = (\hat{f}_1, \dots, \hat{f}_{J-1})'$, whose corresponding covariance matrix is Σ^*/n . The mean of $\hat{\mathbf{f}}^*$ is $P^*\mathbf{g}$, where $P^*_{(J-1) \times K}$ is obtained from P by deleting its last column. Since the distribution of $\hat{\mathbf{f}}^*$ is asymptotically multivariate normal, a solution $\hat{\mathbf{g}}$ to:

$$\min_{\mathbf{g}} (\hat{\mathbf{f}}^* - P^*\mathbf{g})' \Sigma^{*-1} (\hat{\mathbf{f}}^* - P^*\mathbf{g}), \quad (5)$$

$$\text{s.t. } 0 \leq g_i \leq 1, \sum g_i = 1,$$

is asymptotically an mle estimator for \mathbf{g} . Note!, we write ‘an mle’ rather than ‘the mle’ since a solution and an mle are not necessarily unique. See, also Remark 1 below. Practically, Σ^* is replaced by its estimate, which is obtained by utilizing the multinomial distribution of $n\hat{\mathbf{f}}$. A special care should be taken when estimating Σ^{*-1} , since Σ^* might be close to being singular. Our approach in our numerical work was to add 0.001 to the diagonal of the regular estimator of the covariance matrix of a multinomial vector, then take its inverse as the estimator of Σ^{*-1} .

Calibration.

Suppose there is a function A , for which it is known that $E_G A(\theta) = a$. In such a case we may add to the above linear programming the linear constraint:

$$\sum g_k A(s_k) = a.$$

Similarly, when there are a few such functions A_1, \dots, A_b .

The numeric work in this paper was done by applying the quadratic programming function *ipop*, from the R-package *kernelab*, Karatzoglou, et. al. (2004).

Remark 1:

It may be concluded from Lindsay and Roeder (1993) or Lindsay (1995) that when there are only J possible values to Y , there exists an mle for G , that has $J - 1$ points or less in its support. Thus, we can not expect consistency of an arbitrary mle estimator \hat{G} , unless the support of G is known to have no more than $J - 1$ points. However, if Y is obtained by a discretization of a continuous observation, which may become more and more delicate as n grows, we may expect consistency when $J = J_n \rightarrow \infty$.

Furthermore, by adding calibration constraints, we might get an mle which has a larger support and the corresponding estimator \hat{G} is a better approximation of G .

2.2. Deconvolution for estimation of the joint distribution of a latent and an observed variables.

In the previous section we considered the problem of estimating the distribution G of a latent variable θ . In this section we will generalize the method to estimate the joint distribution of a latent variable θ and an observed variable X , where $(X, Y, \theta) \sim G^*$. Let $(X_i, Y_i, \theta_i) \sim G^*$, $i = 1, \dots, n$ be independent. We only observe $\mathcal{T}(X_i, Y_i)$, $i = 1, \dots, n$ for some \mathcal{T} , and the estimation is based only on those observed values.

The variables X_i are discrete, their possible values are x_1, x_2, \dots, x_L .

Our goal is to estimate the joint distribution of θ and X , which is determined by

$$g_{lk} = P_{G^*}(X = x_l, \theta = s_k), \quad l = 1, \dots, L, \quad k = 1, \dots, K,$$

we denote $\mathbf{g} = (g_{11}, g_{12}, \dots, g_{LK}) \equiv (g_1, g_2, \dots, g_{L \times K})$, note the dual indexing of the vector \mathbf{g} .

Let t_1, \dots, t_Q , be the distinct values of $\mathcal{T}(x_l, y_j)$, $l = 1, \dots, L$, $j = 1, \dots, J$. We assume that the conditional distribution $G^*(\mathcal{T} = t | X = x, \theta = s)$ is known, thus the $L \times K$ pairs $\mathbf{v}_{lk} = (x_l, s_k)$, $l = 1, \dots, L$, $k = 1, \dots, K$, play the role of the "parameter" that governs the conditional distribution. Denote by $\mathbf{v} = (\mathbf{v}_{11}, \mathbf{v}_{1,2}, \dots, \mathbf{v}_{LK}) \equiv (\mathbf{v}_1, \dots, \mathbf{v}_{L \times K})$, the vector of "parameters", note the dual indexing of the vector \mathbf{v} .

As in the previous subsection, denote by

$$p_{\mathbf{j}\mathbf{k}} = P(\mathcal{T} = t_{\mathbf{j}} | (X, \theta) = \mathbf{v}_{\mathbf{k}}), \quad \mathbf{j} = 1, \dots, Q, \quad \mathbf{k} = 1, \dots, L \times K,$$

let $P = (p_{\mathbf{j}\mathbf{k}})$ be the corresponding matrix as in the previous subsection.

Given n observations, let $f_{\mathbf{j}} = P(\mathcal{T}(X_i, Y_i) = t_{\mathbf{j}})$, $\mathbf{j} = 1, \dots, Q$, let $\mathbf{f} = (f_1, \dots, f_Q)'$, then

$$\mathbf{f} = P\mathbf{g}.$$

Let $\hat{f}_{\mathbf{j}}$, $\mathbf{j} = 1, \dots, Q$ be the proportion of observations i for which $\mathcal{T}(X_i, Y_i) = t_{\mathbf{j}}$, let $\hat{\mathbf{f}}^* = (\hat{f}_1, \dots, \hat{f}_{Q-1})$. Then $\mathbf{f}^* \equiv E\hat{\mathbf{f}}^* = P^*\mathbf{g}$, for the matrix P^* , which is obtained from P , as in the previous subsection, by deleting its last row.

Let Σ^*/n be the covariance matrix of $\hat{\mathbf{f}}^*$, and suppose it is non-singular. Then,

$$\min_{\mathbf{g}} (\hat{\mathbf{f}}^* - P^*\mathbf{g})' \Sigma^{*-1} (\hat{\mathbf{f}}^* - P^*\mathbf{g}), \quad (6)$$

s.t. $0 \leq g_{lk} \leq 1$, $\sum_{lk} g_{lk} = 1$,
is asymptotically an mle estimator for \mathbf{g} .

Calibration

The above quadratic programming may incorporate various additional linear constraints. For example suppose that there is an indicator $I = I(X)$, $I = 1$ if the corresponding measurement was taken from a male, $I = 0$ otherwise. Suppose it is known that $P_{G^*}(I = 1) = 0.5$. Then the constraint

$$\sum_{(l,k): I(X_l)=1} g_{lk} = 0.5,$$

may be added to the quadratic programming defined in (6).

3. Confidence intervals and linear optimization.

We consider the setup of the previous section, where we observe i.i.d $(X_i, Y_i, \theta_i) \sim G^*$, $i = 1, \dots, n$.

Suppose it is desired to estimate the expectation

$$T = E_{G^*} h(X, \theta) = \sum_{k,l} (x_l, s_k) g_{lk}.$$

Note, a simple modification of the treatment bellow applies also for expectations of the form $T = E_{G^*} h(X, Y, \theta)$; however in order to simplify the notations we consider the above functionals.

It is of an interest to obtain a confidence interval for T , this could reassure that an mle estimator (recall, often the mle is not unique), is giving a reliable estimate.

Let $\hat{\mathbf{f}}^*$ and Σ^* be as in the previous section. Suppose that Σ^* is non-singular. Let $\hat{\Sigma}^*$ be the empirical covariance matrix. Then as the sample size approaches infinity $\hat{\Sigma}^{*-1}$ approaches Σ^{*-1} in probability. Furthermore, the distribution of $\sqrt{n}\hat{\mathbf{f}}^*$ converges weakly to a multivariate normal distribution with covariance matrix Σ^* . Recall, under the general setup of subsection 3.2, we observe $\mathcal{T}(X, Y)$, whose support is of size Q .

Consider the solution of the following problem, of linear optimization under convex constraints.

$$\begin{aligned} T_U &= \max_g \sum_{l,k} h(x_l, s_k) g_{lk} \\ T_L &= \min_g \sum_{l,k} h(x_l, s_k) g_{lk} \end{aligned} \quad (7)$$

s.t.

$$n(\hat{\mathbf{f}}^* - P^* \mathbf{g})' \hat{\Sigma}^{*-1} (\hat{\mathbf{f}}^* - P^* \mathbf{g}) < \chi_{(Q-1), 1-\alpha}^2,$$

$$0 \leq g_{lk} \leq 1, \quad \sum_{l,k} g_{lk} = 1,$$

in the above $\chi_{(Q-1), 1-\alpha}^2$ is the critical value of the appropriate α -level χ^2 test with $Q - 1$ degrees of freedom. As before, additional convex calibration constraints may be added if available.

Theorem 1: If Σ^* is non-singular, then (T_L, T_U) is a $(1 - \alpha)$ level confidence interval, asymptotically as $n \rightarrow \infty$.

The above theorem is for discrete variables θ , X and Y . For continuous cases a discretization should be done. The general guide lines for discretization is that Q will be of size $o(n)$, so that there will be enough observations in each of the $Q - 1$ "cells" and the asymptotic χ_{Q-1}^2 distribution will hold; the considerations involved in the discretization of θ and X have to do with the complexity of the convex optimization. A formal asymptotic treatment of the discretization is beyond the scope of this paper.

4. Non-response and Empirical Bayes type Horvitz Thompson estimators.

A general survey from a population of size N indexed by $\{1, \dots, N\}$, may be described as follows. Each subject i , $i = 1, \dots, N$, in the population is sampled with probability π_i for an interview, but once subject i is sampled a response from that subject is obtained with probability $p_i^* \leq 1$. Let \mathbf{S} be the random set of indices, corresponding to subjects who i) were sampled for an interview ii) responded. Then, for subject i , $P(i \in \mathbf{S}) = \pi_i p_i^* = p_i$, $i = 1, \dots, N$. We define the indicator random variable I_i , $I_i = 1$ iff $i \in \mathbf{S}$; denote $P(I_i = 1) = p_i$.

In many surveys the subjects are equally likely to be sampled to the survey, i.e., $\pi_i \equiv \pi$ are all equal. In the following we treat this case. Thus, w.l.o.g., we may assume that $\pi = 1$, and $p_i = p_i^*$. Modification of the treatment bellow applies when π_i , $i = 1, \dots, N$, may have 'a few' possible values.

We will apply our deconvolution technique and the empirical Bayes ideas, to provide Horvitz Thompson type of estimators in the context of Empirical Bayes.

We model the items of the size N population, as realizations of N , i.i.d random vectors (X_i, Y_i, p_i, I_i) , which are distributed like $(X, Y, p, I) \sim G^*$, X is the variable of interest. The joint distribution of X and p is arbitrary and the conditional distribution of Y conditional on X and p is known;

$$p_i = P(I_i = 1) = P(I_i = 1|X_i, p_i). \quad (8)$$

In order to fix ideas think of X_i as an employment-status of item i , I_i indicator of the event "item i was sampled for a survey and responded". In one of our examples in the sequel, Y_i is the number of attempts until a response was obtained from subject i , where there are at most M_0 attempts. Thus, in this example $I_i = 0$ iff $Y_i > M_0$. We model $Y \sim F_p = \text{Geometric}(\tilde{p})$ for $\tilde{p} = \tilde{p}(p)$, $p = 1 - (1 - \tilde{p})^{M_0}$.

Truncated versus censored observations. We will consider two different setups. In one setup the event $I_i = 0$ means that the observation is truncated, i.e., we do not know about variables with $I_i = 0$, and thus, our available observations may be considered as an i.i.d sample from the distribution, denoted G^{*t} , of (X, Y, p, I) , conditional on $I = 1$. Another setup is of censored observations where we do know about the event $I_i = 0$; e.g., in the example where Y_i is the number of visits until a response, the event $I_i = 0$ implies $Y_i > M_0$. The two setups lead to two versions of our general deconvolution technique, in the truncated setup we estimate the joint distribution of p and X under G^{*t} , while in the censored setup we estimate the joint distribution of p and X under G^* . The joint distribution of X and p under G^* and G^{*t} will be denoted by G and by G^t correspondingly.

4.1. Empirical Bayes type Horvitz Thompson estimators.

Suppose we want to estimate $T = E \sum_{i=1}^N X_i$. We now present three unbiased estimators. Those are in fact pseudo-estimators since they are functions of the unknown p_i , however they will be modified later to become legitimate estimators.

$$\begin{aligned} \hat{T}_0 &= \sum \frac{X_i}{p_i} I_i \equiv \sum X_i I_i A_0^i, \\ \hat{T}_1 &= \sum X_i I_i E\left(\frac{1}{p_i} | X_i, I_i = 1\right) \equiv \sum X_i I_i A_1^i, \\ \hat{T}_2 &= \sum \frac{X_i}{E(p_i | X_i)} I_i \equiv \sum X_i I_i A_2^i. \end{aligned}$$

In the above A_j^i are implicitly defined. The estimator \hat{T}_0 is basically the standard Horvitz Thompson estimator.

Theorem 2:

- i) Under the condition $p_i > 0$ w.p.1, $E(\hat{T}_1) = E(\hat{T}_0) = T$. Under the (weaker) condition $E(p|X) > 0$ w.p.1, $E(\hat{T}_2) = T$.
- ii) Under the condition $p_i > 0$ w.p.1, $Var(\hat{T}_1) \leq Var(\hat{T}_0)$.
- iii) Under the condition $p_i > 0$ w.p.1, $\hat{T}_2 = \hat{T}_1$.

Proof: We prove the theorem for the case $N = 1$.

i) $E\hat{T}_0 = T$ follows immediately, similarly to the implication for a standard Horvitz-Thompson estimator. $E\hat{T}_1 = T$ follows since

$$EE(\hat{T}_0|X, I) = E\hat{T}_1.$$

Assume that $E(p|X) > 0$, w.p.1, then

$$E(\hat{T}_2) = E \frac{XI}{E(p|X)} = EE\left(\frac{XI}{E(p|X)}|X\right) = EX = T,$$

the third equality in the above follows since by (8) $E(I|X) = E(p|X)$.

- ii) The assertion follows by a Rao-Blackwell argument, due to the above conditional expectation representation.
- iii) The assertion follows since

$$dG^*(p|I = 1, X_i = x) = \frac{pdG^*(p|X_i = x)}{\int pdG^*(p|X = x_i)},$$

$$\text{whence } A_1^i = \int \frac{1}{p} dG^*(p|I = 1, X = x_i) = \frac{1}{E(p|X=x_i)} = A_2^i.$$

In practice the terms A_j^i , $j = 1, 2$ are unknown, they will be estimated using our deconvolution method through the estimation of the joint distribution of X and p , under G^* and G^{*t} respectively, for $j = 1, 2$.

For every i , $i = 1, \dots, n$, define

$$\tilde{A}_2^i = 1/E_{\hat{G}}(p|X = X_i);$$

here \hat{G} is the deconvolution estimator for G , the joint distribution of X and Y under G^* .

In the truncated setup the role of G^* in our deconvolution method is played by the conditional distribution G^{*t} . Now,

$$\tilde{A}_1^i = E_{\hat{G}^t}\left(\frac{1}{p}|X = X_i\right);$$

here \hat{G}^t is the estimated joint distribution of X and p under G^{*t} .

We now present the legitimate versions of \hat{T}_1 and \hat{T}_2 , i.e., estimators which are functions only of the available observations,

$$\tilde{T}_1 = \sum X_i I_i \tilde{A}_1^i, \tag{9}$$

$$\tilde{T}_2 = \sum X_i I_i \tilde{A}_2^i. \quad (10)$$

We relate the estimators \tilde{T}_1 and \tilde{T}_2 through the Horvitz-Thompson estimator. However, in fact the estimation of T under the censored setup may be done without the mediation of the Horvitz-Thompson estimator. In fact, an mle estimator for $T = NE(X)$ under the censored setup is:

$$\tilde{T}_3 = NE_{\hat{G}}X = N \sum_l x_l \sum_k \hat{g}_{lk},$$

for a corresponding, mle, estimator $\hat{\mathbf{g}}$. Asymptotically $\tilde{T}_2 \approx \tilde{T}_3$. This may be seen, by the following. Denote by m_l , $l = 1, \dots, L$, the number of indices i satisfying $X_i = x_l$, then $Em_l = NE(p|X = x_l) \sum_k g_{lk}$, whence for large N , $m_l \approx NE_{\hat{G}}(p|X = x_l) \sum_k \hat{g}_{lk}$; note that $\tilde{T}_2 = \sum_l x_l m_l / E_{\hat{G}}(p|X = x_l)$. The later version, \tilde{T}_3 , is better suited compared to \tilde{T}_2 , for deriving a confidence interval for T by the method that is given in Section 3.

There are a few advantages to \tilde{T}_2 compared to \tilde{T}_1 , the obvious one is that it is defined also when the event $p = 0$ has a positive probability. The other advantage is since that in the estimation of A_2 we use some additional censored information, which is not available (i.e., truncated), in the estimation of A_1 . Avoiding possible near singularity for small p , involved in the estimation of A_1^i , is another advantage in attempting to estimate A_2^i when possible. Finally, typically there is an available external information about the distribution G^* , that may be used through calibration, while that information is typically unknown under G^{*t} .

In the following simulation sections, we will apply our estimators in the estimation of the expected proportion α_{x_l} of items with a corresponding $X = x_l$. Their expected total number is estimated by

$$\tilde{T}_{x_l}^j = \sum_{i: X_i = x_l} \tilde{A}_j^i, \quad (11)$$

for $j = 1, 2$, for the truncated and censored setups correspondingly. The following formula applies for the truncated and censored estimators for the proportion α_{x_l} when setting $j = 1, 2$ correspondingly,

$$\hat{\alpha}_{x_{l_0}}^j = \frac{\sum_{i: X_i = x_{l_0}} \tilde{A}_j^i}{\sum_l \sum_{i: X_i = x_l} \tilde{A}_j^i}. \quad (12)$$

In the next section we will use as a benchmark the following estimator, that could be used by an ‘oracle’ that knows p_i $i = 1, \dots, N$. Such an oracle could estimate the size of the population with corresponding $X = x_l$, by $\sum_{i: X_i = x_l} \frac{I_i}{p_i}$. The corresponding oracle estimator for $\alpha_{x_{l_0}}$ would be:

$$\text{oracle}_{x_{l_0}} = \frac{\sum_{i: X_i = x_{l_0}} \frac{I_i}{p_i}}{\sum_l \sum_{i: X_i = x_l} \frac{I_i}{p_i}} \quad (13)$$

5. Simulations

Consider a survey where in its first stage an initial subset of the population is sampled and in the next stage there is an attempt to interview each sampled subject. As mentioned, we assume that each subject in the population is equally likely to be sampled in the first stage, with sampling probability $\pi_i \equiv \pi$, $i = 1, \dots, N$; w.l.o.g $\pi = 1$.

Suppose our policy is to make at most M_0 attempts in order to obtain a response from a sampled subject, however obviously if a response is obtained in the $j < M_0$ attempt, no further attempts are made. We model the number of attempts until a response is obtained by subject i , by a Geometric random variable with a success probability \tilde{p}_i , $i = 1, \dots, N$; assume $0 < \min_i \tilde{p}_i$. Let Y_i denote the number of attempts until a response was obtained.

Assuming $\pi_i \equiv 1$, the probability p_i of subject i , $i = 1, \dots, N$ to be in the set \mathcal{S} , of items that i) were sampled for the survey and ii) responded, is:

$$p_i = 1 - (1 - \tilde{p}_i)^{M_0}. \quad (14)$$

Thus, there is a one to one correspondence between \tilde{p}_i and $p_i = P(I_i = 1) = P(I_i = 1|X_i, p_i) = P(Y_i \leq M_0|X_i, p_i)$.

5.1. Truncated setup

We are interested in the estimation of the joint distribution of X and p under G^{*t} , i.e., conditional upon $I = 1$.

Note, the distribution of Y_i , conditional on $i \in \mathcal{S}$ is:

$$P(Y = j|p) = \frac{(1 - \tilde{p})^{j-1} \tilde{p}}{1 - (1 - \tilde{p})^{M_0}}, \quad j = 1, \dots, M_0; \quad (15)$$

here $\tilde{p} = \tilde{p}(p)$, as given in (14).

Denote the distribution of Y_i , $i \in \mathcal{S}$, given in (15) by F_{p_i} .

Given a grid of points $\{s_1, \dots, s_K\}$ we define the vector $\mathbf{v} = ((x_1, s_1), (x_1, s_2), \dots, (x_L, s_K)) = (\mathbf{v}_1, \dots, \mathbf{v}_{L \times K})$. The possible outcomes, $\mathcal{T}(X, Y) = (X, Y)$ are denoted $((x_1, 1), (x_1, 2), \dots, (x_L, M_0)) = (t_1, \dots, t_Q)$, $Q = L \times M_0$. As in subsection (2.2), let $\mathbf{f} = (f_1, \dots, f_Q)$, where $f_{\mathbf{j}} = P(\mathcal{T}(X, Y)) = t_{\mathbf{j}}$, $\mathbf{j} = 1, \dots, Q$. For every \mathbf{j} , $\mathbf{j} = 1, \dots, Q$, we denote $t_{\mathbf{j}} = (t_{j1}, t_{j2})$, for every \mathbf{k} , $\mathbf{k} = 1, \dots, K \times L$, $\mathbf{v}_{\mathbf{k}} = (v_{k1}, v_{k2})$.

Denote

$$p_{\mathbf{j}\mathbf{k}} = P(\mathcal{T} = t_{\mathbf{j}}|(X, p) = \mathbf{v}_{\mathbf{k}}), \quad \mathbf{j} = 1, \dots, Q, \quad \mathbf{k} = 1, \dots, L \times K.$$

This defines the matrix $P = (p_{\mathbf{j}\mathbf{k}})$ as explained in the previous section,

$$p_{\mathbf{j}\mathbf{k}} = \begin{cases} 0 & v_{k1} \neq t_{j1} \\ P(Y = t_{j2}|p = v_{k2}) & v_{k1} = t_{k1}. \end{cases}$$

Note, $\mathbf{f} = P\mathbf{g}$. We proceed as in subsection (2.2) to derive an estimator for \mathbf{g} . In turn we obtain the estimators $\hat{a}_{x_l}^1$, $l = 1, \dots, L$, as in (12).

5.2. Censored setup.

We will repeat the estimation of α_{x_l} , $x_l = 0, 1$, under the same setup, estimating the proportions by $\hat{\alpha}_{x_l}^2$, i.e, the censored version of (12) .

In the current setup we observe $\mathcal{T}(X_i, Y_i)$, where

$$\mathcal{T}(X_i, Y_i) = \begin{cases} (X_i, Y_i) & Y_i \leq M_0 \\ NR & Y_i > M_0 \end{cases}$$

Here "NR" abbreviate "Non-Response" and the outcome NR implies that $Y_i > M_0$. We denote the possible outcomes by $(t_1, \dots, t_Q) = ((X_1, 1), (X_1, 2), \dots, (X_L, M_0), "NR")$. The number of possible values of \mathcal{T} is $Q = (M_0 \times L) + 1$. As in subsection (2.2), let \mathbf{f}^* be the vector of expected proportions of the $Q - 1$ possible outcomes when excluding the outcome "NR"; let \mathbf{v} be the $L \times K$ dimensional vector, as in (2.2).

We write $\mathbf{v}_{\mathbf{k}} \equiv (v_{\mathbf{k}1}, v_{\mathbf{k}2})$, $t_{\mathbf{j}} \equiv (t_{\mathbf{j}1}, t_{\mathbf{j}2})$.

$$p_{\mathbf{j}\mathbf{k}} = \begin{cases} 0 & v_{\mathbf{k}1} \neq t_{\mathbf{j}1} \\ (1 - \tilde{p}_{\mathbf{k}})^{t_{\mathbf{j}2}} \tilde{p}_{\mathbf{k}} & v_{\mathbf{k}1} = t_{\mathbf{j}1}. \end{cases}$$

Here $\tilde{p}_{\mathbf{k}} = \tilde{p}_{\mathbf{k}}(v_{\mathbf{k}2})$, is the probability of success of the Geometric random variable Y , while $v_{\mathbf{k}2}$ is the probability of success within M_0 trials, $v_{\mathbf{k}2} = 1 - (1 - \tilde{p}_{\mathbf{k}})^{M_0}$, as explained in the previous subsection.

We proceed as in subsection (2.2), to obtain the deconvolution estimator for $\mathbf{g} = (g_{11}, g_{12}, \dots, g_{LK})$ that determines G , the joint distribution of X and p under G^* . The estimated \hat{G} defines the estimator $\hat{\alpha}_{x_l}^2$, $l = 1, \dots, L$, for α_{x_l} , as explained in (12).

5.3. Numerical experiments

In the following we simulate populations of size N , where we randomly assigned to $N^0 \sim \text{Binomial}(N, 0.5)$ items a corresponding value $X = 0$, and to the remaining N^1 a corresponding $X = 1$ was assigned. A value \tilde{p} , of a response probability in a single attempt, was randomly assigned to each of the N^0 items independently, under a distribution \tilde{G}_0 . Similarly a value \tilde{p} was assigned randomly to each of the N^1 items based on a distribution \tilde{G}_1 . A corresponding pair (G_0, G_1) , of distributions of the possible values of response probabilities p is determined. Let $\alpha_{x_l} = 0.5$, $x_l = 0, 1$, be the expected proportion of items with a corresponding $X = x_l$. Finally, for each item i , $i = 1, \dots, N$, a Geometric random variable $Y_i \sim \text{Geometric}(\tilde{p}_i)$ was simulated.

We simulated scenarios with $N = 1000$ and $N = 10000$. The cases $M_0 = 4, 6, 8$, were studied for each of the following three classes of pairs of distributions $(\tilde{G}_0, \tilde{G}_1)$, parametrized by γ .

Two Points. The distribution \tilde{G}_0 has a two points support, at the points 0.5 and 0.9, with probability mass 0.5 at each.

The distribution $\tilde{G}_1 \equiv \tilde{G}_1^\gamma$ is a $(-\gamma)$ translation of \tilde{G}_0 . We present results for the cases $\gamma = 0.1, 0.2, 0.3, 0.4$.

Uniform. The distribution \tilde{G}_0 is uniform on the interval $(0.1, 1)$. The distribution $\tilde{G}_1 \equiv \tilde{G}_1^\gamma$, is a mixture of \tilde{G}_0 and a point mass at 0.1, where the mixing weights are $(1 - \gamma)$ and γ correspondingly. We present results for $\gamma = 0.1, 0.2, 0.3, 0.4$.

Normal. The distribution \tilde{G}_0 is a $N(0.5, 0.1)$, ‘rounded up’ to 0.1 and ‘rounded down’ to 1. The distribution $\tilde{G}_1 \equiv \tilde{G}_1^\gamma$ is $N(0.5 - \gamma, 0.1)$ ‘rounded up’ to 0.1 and ‘rounded down’ to 1. We present results for $\gamma = 0.1, 0.2, 0.3, 0.4$.

In the simulations we compared the performance of the following estimators for $\alpha_0 = 0.5$. The naive estimator, that estimates α_0 by the sample proportion, i.e., the proportion among responders, of items i with $X_i = 0$; the estimators $\hat{\alpha}_0^j$, $j = 1, 2$, that correspond to the truncated and censored setups, as given in (12); the ‘oracle’ estimator as in (13).

The grid points $\{s_1, \dots, s_K\}$ taken as the support of G , are induced by $\{\tilde{s}_1 = 0.1, \tilde{s}_2 = 0.12, \dots, \tilde{s}_K = 1\}$ that were taken as the support of \tilde{G} .

The following two tables correspond to the cases $N = 1000$ and $N = 10000$. The columns S-naive, S- $\hat{\alpha}_0^1$, S- $\hat{\alpha}_0^2$, S-oracle correspond to the square root of the simulated mean squared error of each of the corresponding methods, based on 1000 repetitions. The columns m-naive, m- $\hat{\alpha}_0^1$, m- $\hat{\alpha}_0^2$, correspond to the simulated average of each of the corresponding methods; the simulated mean of the oracle’s estimator was virtually 0.5 and thus not presented.

It may be seen that $\hat{\alpha}_0^2$, clearly dominates $\hat{\alpha}_0^1$, as may be expected. The performance of all the methods is improved by an increase in M_0 , but the improvement is much sharper for $\hat{\alpha}_0^j$, $j = 1, 2$. For $\gamma = 0$, the setup would become Missing at Random, in which the naive estimator is the best. As γ increases, the other methods dominate the naive. It may be seen that for large enough M_0 and γ , in all of our simulated configurations $\hat{\alpha}_0^2$, dominates the naive. In the uniform case when $N = 1000$, $\hat{\alpha}_0^1$ does not dominate the naive in any of the configurations, however when we let $M_0 = 10$, we get domination of $\hat{\alpha}_0^1$, specifically, for $N = 1000$, $M_0 = 10$, $\gamma = 0.4$, S- $\hat{\alpha}_0^1 = 0.0279$, compared to S-naive = 0.0383.

The performance of the estimator $\hat{\alpha}_0^2$, is comparable to that of the oracle when $M_0 = 8$, and it is amazingly close to it in the Two-Points case. It may be seen that our methods reduce the bias. This is important beyond the reduction of the mse, since often the estimators arrive as a time-series and the final estimators involve additional smoothing of the time-series. Obviously, smoothing around the true value gives further reduction in mse, compared to smoothing of a biased sequence.

Finally, an important ‘moral’ from the two tables is that an increase in M_0 is much more important for risk reduction, relative to an increase in the sample size. For example, in the setup of Two-Points, $\gamma = 0.4$, $M_0 = 6$, $N = 1000$, we have S- $\hat{\alpha}_0^1 = 0.0274$, while for $\gamma = 0.4$, $M_0 = 4$, $N = 10000$, the mse is increased and S- $\hat{\alpha}_0^1 = 0.0478$. Obviously, the number of interviewing attempts in the first case is smaller than that in the second case.

TABLE 1
 $N=1000$

\tilde{G}	M_0	γ	m-naive	m- $\hat{\alpha}_0^1$	m- $\hat{\alpha}_0^2$	S-naive	S-oracle	S- $\hat{\alpha}_0^1$	S- $\hat{\alpha}_0^2$
TwoPts	4	0.1	0.4909	0.4206	0.4963	0.0184	0.0161	0.0868	0.0186
TwoPts	4	0.2	0.4743	0.4094	0.4891	0.0304	0.0165	0.0996	0.0250
TwoPts	4	0.3	0.4470	0.3867	0.4766	0.0556	0.0173	0.1230	0.0405
TwoPts	4	0.4	0.3978	0.3456	0.4532	0.1035	0.0192	0.1618	0.0668
TwoPts	6	0.1	0.4966	0.4815	0.4995	0.0164	0.0161	0.0300	0.0164
TwoPts	6	0.2	0.4872	0.4823	0.4978	0.0208	0.0165	0.0335	0.0173
TwoPts	6	0.3	0.4663	0.4726	0.4937	0.0373	0.0162	0.0417	0.0203
TwoPts	6	0.4	0.4221	0.4358	0.4846	0.0796	0.0178	0.0731	0.0274
TwoPts	8	0.1	0.4978	0.4975	0.4992	0.0156	0.0154	0.0172	0.0155
TwoPts	8	0.2	0.4933	0.5007	0.4996	0.0173	0.0160	0.0200	0.0160
TwoPts	8	0.3	0.4788	0.5022	0.4990	0.0268	0.0165	0.0245	0.0169
TwoPts	8	0.4	0.4394	0.4762	0.4948	0.0629	0.0178	0.0349	0.0185
Uniform	4	0.1	0.4855	0.3739	0.4921	0.0224	0.0181	0.1335	0.0446
Uniform	4	0.2	0.4682	0.3638	0.4816	0.0360	0.0184	0.1435	0.0548
Uniform	4	0.3	0.4504	0.3562	0.4777	0.0530	0.0201	0.1516	0.0609
Uniform	4	0.4	0.4301	0.3509	0.4710	0.0720	0.0197	0.1571	0.0664
Uniform	6	0.1	0.4882	0.4441	0.4952	0.0205	0.0174	0.0629	0.0287
Uniform	6	0.2	0.4738	0.4399	0.4893	0.0312	0.0174	0.0679	0.0340
Uniform	6	0.3	0.4597	0.4347	0.4860	0.0438	0.0176	0.0735	0.0371
Uniform	6	0.4	0.4457	0.4314	0.4858	0.0570	0.0183	0.0770	0.0388
Uniform	8	0.1	0.4908	0.4757	0.4973	0.0189	0.0166	0.0340	0.0224
Uniform	8	0.2	0.4794	0.4709	0.4941	0.0261	0.0162	0.0373	0.0238
Uniform	8	0.3	0.4679	0.4687	0.4937	0.0362	0.0172	0.0408	0.0256
Uniform	8	0.4	0.4555	0.4634	0.4913	0.0476	0.0173	0.0449	0.0255
Normal	4	0.1	0.4792	0.3570	0.4966	0.0267	0.0168	0.1492	0.0227
Normal	4	0.2	0.4422	0.3485	0.4917	0.0602	0.0176	0.1594	0.0295
Normal	4	0.3	0.3859	0.3414	0.4863	0.1156	0.0199	0.1679	0.0404
Normal	4	0.4	0.3231	0.3332	0.4833	0.1778	0.0211	0.1738	0.0471
Normal	6	0.1	0.4902	0.4571	0.4989	0.0195	0.0169	0.0523	0.0184
Normal	6	0.2	0.4664	0.4489	0.4955	0.0375	0.0169	0.0631	0.0214
Normal	6	0.3	0.4223	0.4380	0.4920	0.0796	0.0180	0.0744	0.0257
Normal	6	0.4	0.3691	0.4333	0.4919	0.1321	0.0191	0.0782	0.0272
Normal	8	0.1	0.4945	0.4899	0.4987	0.0169	0.0160	0.0232	0.0168
Normal	8	0.2	0.4777	0.4875	0.4968	0.0277	0.0166	0.0277	0.0178
Normal	8	0.3	0.4461	0.4813	0.4964	0.0564	0.0170	0.0350	0.0187
Normal	8	0.4	0.4016	0.4762	0.4962	0.0999	0.0183	0.0401	0.0196

TABLE 2
 $N = 10000$

\tilde{G}	M_0	γ	m-naive	m- $\hat{\alpha}_0^1$	m- $\hat{\alpha}_0^2$	S-naive	S-oracle	S- $\hat{\alpha}_0^1$	S- $\hat{\alpha}_0^2$
TwoPts	4	0.1	0.4907	0.4191	0.4974	0.0106	0.0052	0.0837	0.0081
TwoPts	4	0.2	0.4747	0.4119	0.4931	0.0258	0.0054	0.0925	0.0136
TwoPts	4	0.3	0.4469	0.3939	0.4849	0.0534	0.0053	0.1110	0.0261
TwoPts	4	0.4	0.3983	0.3478	0.4648	0.1019	0.0061	0.1559	0.0478
TwoPts	6	0.1	0.4957	0.4786	0.4993	0.0067	0.0051	0.0237	0.0053
TwoPts	6	0.2	0.4870	0.4786	0.4990	0.0139	0.0050	0.0261	0.0058
TwoPts	6	0.3	0.4663	0.4767	0.4976	0.0341	0.0053	0.0303	0.0071
TwoPts	6	0.4	0.4228	0.4459	0.4928	0.0774	0.0057	0.0574	0.0117
TwoPts	8	0.1	0.4985	0.4979	0.5001	0.0052	0.0049	0.0066	0.0050
TwoPts	8	0.2	0.4933	0.4976	0.4999	0.0084	0.0050	0.0088	0.0051
TwoPts	8	0.3	0.4785	0.4999	0.4995	0.0221	0.0053	0.0126	0.0054
TwoPts	8	0.4	0.4395	0.4829	0.4978	0.0607	0.0055	0.0206	0.0060
Uniform	4	0.1	0.4845	0.3666	0.4926	0.0164	0.0057	0.1361	0.0321
Uniform	4	0.2	0.4679	0.3617	0.4852	0.0326	0.0060	0.1413	0.0393
Uniform	4	0.3	0.4504	0.3580	0.4833	0.0499	0.0062	0.1457	0.0437
Uniform	4	0.4	0.4317	0.3524	0.4807	0.0686	0.0061	0.1516	0.0493
Uniform	6	0.1	0.4874	0.4412	0.4945	0.0136	0.0052	0.0612	0.0193
Uniform	6	0.2	0.4741	0.4363	0.4919	0.0264	0.0054	0.0663	0.0224
Uniform	6	0.3	0.4600	0.4338	0.4910	0.0404	0.0057	0.0693	0.0238
Uniform	6	0.4	0.4453	0.4299	0.4891	0.0550	0.0057	0.0736	0.0257
Uniform	8	0.1	0.4897	0.4724	0.4976	0.0116	0.0054	0.0296	0.0119
Uniform	8	0.2	0.4791	0.4692	0.4957	0.0216	0.0053	0.0331	0.0134
Uniform	8	0.3	0.4681	0.4677	0.4958	0.0323	0.0054	0.0349	0.0140
Uniform	8	0.4	0.4563	0.4653	0.4947	0.0440	0.0057	0.0375	0.0144
Normal	4	0.1	0.4792	0.3498	0.4963	0.0215	0.0052	0.1519	0.0124
Normal	4	0.2	0.4432	0.3398	0.4913	0.0571	0.0056	0.1633	0.0196
Normal	4	0.3	0.3863	0.3287	0.4852	0.1138	0.0060	0.1749	0.0296
Normal	4	0.4	0.3223	0.3276	0.4870	0.1777	0.0064	0.1751	0.0297
Normal	6	0.1	0.4892	0.4553	0.4981	0.0120	0.0052	0.0466	0.0076
Normal	6	0.2	0.4656	0.4456	0.4943	0.0348	0.0052	0.0574	0.0119
Normal	6	0.3	0.4224	0.4330	0.4915	0.0778	0.0055	0.0710	0.0155
Normal	6	0.4	0.3695	0.4275	0.4920	0.1306	0.0059	0.0756	0.0148
Normal	8	0.1	0.4942	0.4897	0.4990	0.0077	0.0051	0.0131	0.0058
Normal	8	0.2	0.4786	0.4851	0.4973	0.0221	0.0052	0.0186	0.0072
Normal	8	0.3	0.4460	0.4791	0.4965	0.0543	0.0055	0.0260	0.0080
Normal	8	0.4	0.4030	0.4743	0.4964	0.0971	0.0056	0.0300	0.0081

6. Analysis of real data of Labor Force Survey.

In this section we will apply our method on a real data set from the Labor Force Survey, that is conducted by the Israel Central Bureau of Statistics. The sampling method is 4-8-4 rotating panels, however for our analysis, it may be equivalently treated and described as a 4-in rotation, which is described in the following.

The survey is given to four panels, where each panel is investigated for four consecutive months. Each month one panel finishes its fourth investigation and in the next month it will be replaced by a new panel that will remain for four months. The main purpose of the survey is to estimate the proportion of ‘Unemployment’, ‘Employment’, and those who are ‘Not in Working Force (NWF)’, the last category is of those who do not have a job nor they are looking for one. Denote the corresponding values of our variable X -‘working status’, by 0, 1, 2. We are interested in estimating α_0, α_1 and α_2 . The population of interest is of residents whose age is above 15, and the proportions are with respect to that population. The probability π to be included in the sample is the same for each person. As explained, for our purpose of estimating proportions we assume w.l.o.g that $\pi = 1$.

Temporarily assume that, we have only the data from the panel that is investigated for the fourth time (‘fourth panel’). Its size is about 5000, however, only n responses were obtained, m_l responses from people with working status x_l , $x_l = 0, 1, 2$. The general response rate is about 80 percent in each month. For each of the responding n units there is a corresponding random variable, denoted Y , that counts the number of responses, in the four interviewing attempts. Those with 0 responses are truncated. Indeed the records for the reason of 0 responses were not accurate, and thus we preferred to ignore/truncate the records that correspond to 0 responses. We model the distribution of an observed random variable, i.e. conditional on $i \in \mathbf{S}$ by

$$Y = 1 + W; \quad W \sim \text{Binomial}(3, p).$$

The above model amounts to assuming that the probability of response of unit i , is p_i in all of its four investigation attempts, and responses in different months are independent. Given a grid s_1, \dots, s_k , for the support of the possible values of p , a matrix $P = (p_{jk})$ is defined where $p_{jk} = P(Y = j | p = s_k) = P_{s_k}(1 + W = j)$, $j = 1, 2, 3, 4$, for $W \sim B(3, s_k)$. In our analysis we took the grid 0.1, 0.11, 0.12, ..., 1. The above induces a matrix P^* in a manner similar to the previous sections.

Now, α_{x_l} , $l = 0, 1, 2$, may be estimated by $\hat{\alpha}_{x_l}^1$ as given in (12) for a truncation setup. However, so far we considered only the data from the panel that has four investigations. Indeed the panels that have less investigations will yield poor estimates of $E(1/p | X = x_l, I = 1)$. Our approach is the following hybrid method. We estimate $E(1/p | X = x_l, I = 1)$, $l = 0, 1, 2$, based not only on the data from the current ‘fourth panel’, but, in addition we use the data obtained in the four investigation of the three more panels that had their fourth investigation

in the previous month, two months ago, and three months ago, altogether four panels. Let m_l , be the number of items in the *currently* investigated four panels, with corresponding $X = x_l$, $x_l = 0, 1, 2$. Our hybrid approach is to inflate m_l , which is based on the currently investigated four panels, using the estimated $E(1/p|X = x_l, I = 1)$, $x_l = 0, 1, 2$, which are in turn based on the current as well as ‘historical’ complementary information. The underlying assumption is that $E(1/p|X = x_l, I = 1)$, changes slowly in time and thus, estimating it based on a complementary older data, we still get at least some bias correction. We proceed by estimating by \hat{G}^t , the joint distribution of (X, p) under truncation. Finally, we get the estimator

$$\hat{\alpha}_{x_{l_0}} = \frac{m_{l_0} E_{\hat{G}^t}(1/p|X = x_{l_0})}{\sum_l m_l E_{\hat{G}^t}(1/p|X = x_l)}.$$

Since the true proportions of the various working statuses are unknown, we will first demonstrate the performance of the above estimation method in estimating the following *known* true proportions, based on the responses in a given month.

In one case we estimate the proportion of males in the population, which is known to be 0.4853; their proportion in the survey among responders is about one percent lower. In the other example we estimate the proportion of the group age 20-39. Their known proportion is 0.397 while their, response rate is particularly low, their proportion among the responders is nearly 3 percent lower than their proportion in the population.

Each of the following tables 3 and 4 has three lines that correspond to the data obtained in Aug/2012, Dec/2012, and April/2013. We took periods that are four months apart in order not to have overlapping panels. The general picture persist in other months.

The columns True, Naive, and $\hat{\alpha}$, correspond to the true populations proportion, the sample proportion among responders, and our estimator $\hat{\alpha}$. In each case one may see that $\hat{\alpha}$ corrects the sample proportion in the right direction.

After gaining some confidence in $\hat{\alpha}$, we will now examine its estimates in the estimation of the proportion of ‘Unemployed’, ‘Employed’ and those ‘Not in Working Force’ (NWF). In the following Table 5 the columns Naive and $\hat{\alpha}$ are as before. The column Bureau gives the estimates of the Israel, Central Bureau of Statistics, for the three categories of working statuses. The three parts of the table refer to the three working statuses. The three lines in each part refer to the three months as described before. The Bureau and the $\hat{\alpha}$ estimators ‘correct’ the naive estimator for Employment and NWF, in opposite directions (the official Bureau estimator involves additional seasonal adjustment that we neglect). The estimator of the bureau is obtained through a method that involves calibration in a ‘post-stratification manner’. It seems that the correction of the bureau, of ‘Employment’ and the ‘NWF’ is in the wrong direction. This is indicated also

when imputing missing values based on their values in months where a response was obtained looking also ‘into the future’. On the other hand both the Bureau and $\hat{\alpha}$ correct the unemployment naive estimate by increasing it. This direction of correction of unemployment, is ,again, supported also by an analysis that involves imputation.

TABLE 3
Comparison of estimates of male’s proportion.

	True	Naive	$\hat{\alpha}$
Male	0.4853	0.4752	0.4822
	0.4853	0.4751	0.4819
	0.4853	0.4776	0.4842

TABLE 4
Comparison of estimates of proportion of 20-39 age group.

	True	Naive	$\hat{\alpha}$
Age 20-39	0.3970	0.3664	0.3815
	0.3970	0.3631	0.3984
	0.3970	0.3598	0.3842

TABLE 5
Comparison of unemployment estimates.

	Bureau	Naive	$\hat{\alpha}$
Emp	0.6104	0.5931	0.5761
	0.6081	0.5992	0.5910
	0.6089	0.5986	0.5881
NWF	0.3416	0.3594	0.3748
	0.3465	0.3576	0.3605
	0.3491	0.3621	0.3720
UnEmp	0.0479	0.0475	0.0492
	0.0454	0.0431	0.0484
	0.0420	0.0392	0.0399

References

- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *JRSSB* **57** No.1, 289-300.
- Brown, L.D. and Greenshtein, E. (2009). Non parametric empirical Bayes and compound decision approaches to estimation of high dimensional vector of normal means. *Ann. Stat.* **37**, No. 4, 1685-1704.
- Brown L.D., Greenshtein, E. and Ritov, Y. (2013). The Poisson compound decision revisited. *JASA*. **108** 741-749.
- Efron, B. (2013). Empirical Bayes modeling, computation and accuracy. Manuscript.
- Greenshtein, E., Park, J., and Ritov, Y. (2008). Estimating the mean of high valued observations in high dimensions. *JSTP* **2** No. 3 407-418.
- Koenker, R. and Mizera, I. (2013). Convex optimization, shape constraints, compound decisions and empirical Bayes rules. Manuscript.
- Lee, M., Hall, P., Haipeng, S., Marron, J.S., and Tolle, J. (2013). Deconvolution estimation of mixture distributions with boundaries. (2013). *Electronic J. of Stat.* **7** 323-341.
- Lindsay, B. G. (1995). Mixture Models: Theory, Geometry and Applications. Hayward, CA, IMS.
- Lindsay, B. G. and Roeder, K. (1993). Uniqueness of estimation and identifiability in mixture models. *Canadian Journal of Stat.* **21**, No. 2, 139-147.
- Little, R.J.A and Rubin, D.B. (2002). Statistical Analysis with Missing Data. New York: Wiley.
- Karatzoglou, A., Smola, A., Hornik, K., and Zeileis, A., (2004). An S4 package for kernel methods in R. *Journal of Statistical Software* **11**, No. 9, 1-20.
- Kiefer and Wolfowitz (1956). Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *Ann.Math.Stat.* **27** No. 4, 887-906.
- Sharon L. Lohr (1999). *Sampling Design and Analysis*. Brooks/Cole publishing company.
- Zhang, C-H. (2005). Estimation of sums of random variables: Examples and information bounds. *Ann. Stat.* **33** No.5. 2022-2041.